Managing Usability and Reliability Aspects in Cloud Computing

Maria Spichkova[1], Heinz W. Schmidt[2], Ian E. Thomas[2], Iman I. Yusuf[2]
Steve Androulakis[3] and Grischa R. Meyer[3]
[1]*School of Science, RMIT University, 414-418 Swanston Street, 3001, Melbourne, Australia*
[2]*eResearch, RMIT University, 17-23 Lygon Street, 3053, Carlton, Australia*
[3]*eResearch, Monash University, Wellington Rd, 3800, Clayton, Australia*
{*maria.spichkova, heinz.schmidt, ian.edward.thomas, iman.yusuf*}*@rmit.edu.au,*
{*steve.androulakis, grischa.meyer*}*@monash.edu*

arXiv:1612.01675v1 [cs.SE] 6 Dec 2016

Keywords: Usability, Reliability, Cloud Computing, Modelling, Visualisation

Abstract: Cloud computing provides a great opportunity for scientists, as it enables large-scale experiments that cannot are too long to run on local desktop machines. Cloud-based computations can be highly parallel, long running and data-intensive, which is desirable for many kinds of scientific experiments. However, to unlock this power, we need a user-friendly interface and an easy-to-use methodology for conducting these experiments. For this reason, we introduce here a formal model of a cloud-based platform and the corresponding open-source implementation. The proposed solution allows to conduct experiments without having a deep technical understanding of cloud-computing, HPC, fault tolerance, or data management in order to leverage the benefits of cloud computing. In the current version, we have focused on biophysics and structural chemistry experiments, based on the analysis of big data from synchrotrons and atomic force microscopy. The domain experts noted the time savings for computing and data management, as well as user-friendly interface. [a]

[a]Preprint. Accepted to the 11th International Conference on Evaluation of Novel Approaches to Software Engineering (ENASE 2016). Final version published by SCITEPRESS, http://www.scitepress.org

1 INTRODUCTION

Scientific experiments can be very challenging from a domain point of view, even in the case the computation can be done on a local desktop machine. Instruments such as synchrotrons and atomic force microscopy produce massive amounts of data. Methods used by scientists are frequently implemented in software. The combination of big data and complex computational methods inevitably requires long running computations and demand for high-performance computing (HPC) using cluster or cloud computing (Armbrust et al., 2010; Buyya et al., 2009). HPC and cloud computing marshal large storage resources flexibly and divide tasks and data up to huge numbers of compute cores. For most users, both present new technologies, either computationally and in data management, and both require learning non-standard data management, programming languages and libraries.

In the case of cloud computing, the users have to learn how to work within a cloud-based environment, e.g., how to create and set up virtual machines (VMs), how to collect the results of their experiments, and finally destroy the VMs, etc. Thus, to unlock all the capabilities of cloud computing the science users have to obtain a new set of skills (e.g., knowledge of fault tolerance), which might distract from focusing on the domain specific problems of the experiment.

Cloud computing provides many benefits, e.g., access to online storage and computing resources at a moment's notice. Nevertheless, failure while setting up a cloud-based execution environment or during the execution itself is arguably inevitable: some or all of the requested VMs may not be successfully created/instantiated, or the communication with an existing VM may fail due to long-distance network failure – given clouds data centres are typically remote and communication crosses many network boundaries. Also, one has to realise that all tasks of such parallel computations are required to complete, therefore the failure of any one of them may corrupt the result in some way. Statistically this means that the reliability of the overall task completion is the product of that of the individual tasks – and with very many thou-

sands or millions of compute tasks this may quickly become a vanishable number.

For these reasons, we propose a user-friendly open-source platform that would hide the above problems from the user by incapsulating them in the platform's functionality. The feasibility of our platform is shown using case studies across Theoretical Chemical and Quantum Physics group at the RMIT university. This paper presents a formal model of a cloud-based platform for scalable and fault-tolerant cloud computations as well as its implementation. We focus on scientific computations, i.e., we assume that the users of the platform would be researchers working in the fields of physics, chemistry, biology, etc.

Our solution enables researchers to focus on domain-specific problems, and to delegate to the tool to deal with the detail that comes with accessing high-performance and cloud computing infrastructure, and the data management challenges it poses. Moreover, the platform implements various fault tolerance strategies to prevent the failed execution from causing a system-wide failure, as well as to recover a failed execution.

Outline: The rest of the paper is structured as follows. Section 2 presents core features of the proposed model and its implementation as an open-source cloud-based platform, including the reliability aspects. In Section 3, we discuss the usability features of the platform and how they are reflected in the conducted case studies. Section 4 overviews the related work. Section 5 concludes the paper by highlighting the main contributions, and introduces the future work directions.

2 MODEL OF A CLOUD-BASED PLATFORM

The proposed platform provides access to a distributed computing infrastructure. On the logical level it is modelled as a dynamically built set of *Smart Connectors* (SCs), which handle the provision of cloud-based infrastructure. SCs vary from each other by the type of computation to be supported and/or the specific computing infrastructure to be provisioned.

An SC interacts with a cloud service (Infrastructure-as-a-Service) on behalf of the user. Figure 1 presents the corresponding workflow. With respect to the execution environment, the only information that is expected from the user is to specify the number of computing resources she wishes to use, credentials to access those resources, and the location for transferring the output of the computation. Thus, the user does not need to know about how the execution environment is set up (i.e., how VMs are created and configured for the upcoming simulation), how a simulation is executed, how the final output is transferred and how the environment is cleaned up after the computation completion (i.e., how the VMs are destroyed).

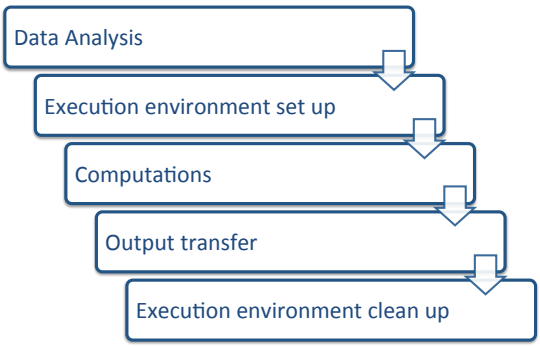

Figure 1: Cloud service: Workflow

Figure 2 shows the logical architecture of an SC. Each SC consist of five logical components, which can also be seen as processes within the framework workflow, presented on Figure 1. The model is based on our previous research on process analysis and specification (Spichkova and Schmidt, 2015; Spichkova, 2011). While developing the model we focused on the understandability and readability aspects (Spichkova, 2013a; Spichkova et al., 2013).

We specify for any process P its entry and exit points by $Entry(P)$ and $Exit(P)$ respectively, and represent a process P (elementary or composed) by the corresponding component specification $PComp$. All the control channels (representing entry and exit points of a process) are drawn as orange dashed lines, the corresponding auxiliary components over these channels are also drawn in orange.

An execution of an SC is called a *job*. An SC executes a user requested *process cP*, which consists of tasks $Task_1, \ldots, Task_{NT}$, which could be executed in iterative manner. In the simplest case, cP consists of a single task that should be executed once only.

A concrete SC is build from a general template by configuration its parameters:

- *DataConstraints* specifies constraints on the user provided input *dataInput*;
- *ExecParamVM* specifies parameters of the job, e.g., which compilers should be installed on the generated VMs;
- *ExecParamT* is a list of the task execution parameters $ExecParamT_1, \ldots, ExecParamT_{NT}$. These parameters specify for each task which data are

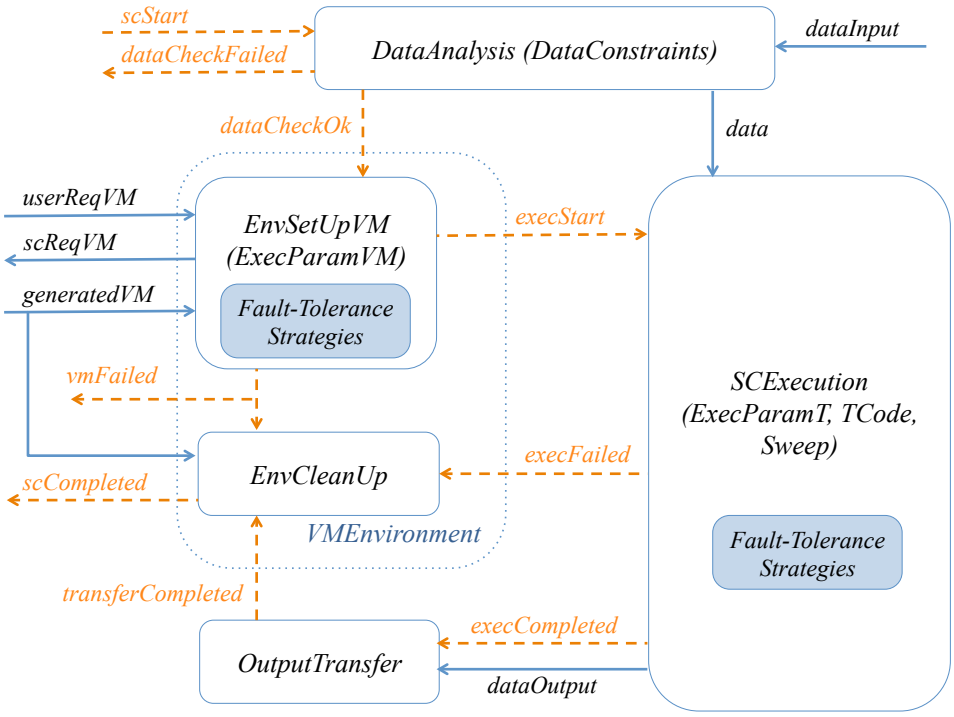

Figure 2: Logical architecture of a *Smart Connector*.

required for its execution, what is a convergence criterion and whether there is any for that task, which scheduling constrains are required, etc.;

- *TCode* presents an actual executable code for the corresponding tasks, in general case it consists of *NT* elements.
- *Sweep* is a list of values to sweep over: With respect to configuring and executing the simulation, the user may set the value and/or ranges of domain specific parameters, and subsequently automatically creating and executing multiple instances of the given SC to sweep across ranges of values.

The first three parameters can be partially derived from *TCode* on the development stage for a concrete Smart Connector.

The *DataAnalysis* component is responsible for the preliminary check whether the user *dataInput* satisfies the corresponding *DataConstraints*, both on syntactical and on semantical level. The *DataAnalysis* process is started by receiving an *scStart* signal from the user. If the data check was successful, the *VMEnv* component is activated by signal *dataCheckOk* and the data are forwarded to the *SCExecution* component, otherwise the process is stopped and the user receives an error message *dataCheckFail*.

The *EnvSetUpVM* component is responsible for the communication with the cloud to obtain a number of VMs that is enough for the task according to the user requests *userReqVM*. The user request *userReqVM* is a pair of numbers (iN, mN), where iN is an ideal and mN is a minimal (from the user's point of view) number of VMs required for the experiment. The *EnvSetUpVM* component requests from the cloud iN VMs.

Fault-Tolerance properties of *EnvSetUpVM*:
If some of the requested VMs are not created/instantiated successfully (i.e. only j VMs are successfully created, where $0 \leq j < mN$. The SC will employ various strategies to create the remaining VMs: it will retry to make either a block request to create $mN - j$ VMs at once or a single request at a time. For these purposes, one of the parameters within *ExecParamVM* have to be *RetryLimit*, which limits the number of retries.

The number of VMs generated by cloud (the list *generatedVM* represents the list of the corresponding identifiers, e.g., IP addresses) has to fulfil the following property

$$mN \leq length(generatedVM) \leq iN \qquad (1)$$

If Equation 1 is not fulfilled (more precisely, if $length(generatedVM) < mN$, because the cloud never provides a number of VMs larger than requested), or the bootstrapping failed (e.g. due lost connection to

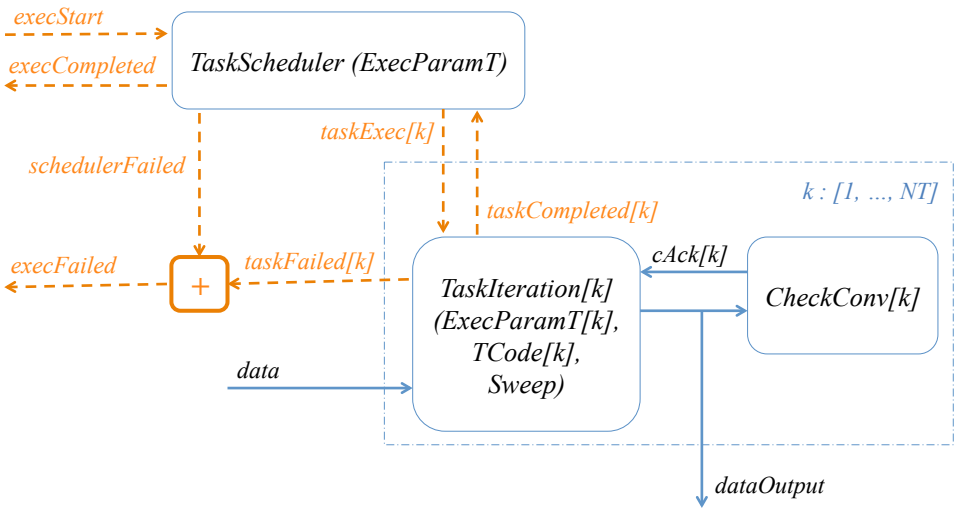

Figure 3: *SCExecution* subcomponent of Smart Connector.

the cloud) the process is stopped and the user receives an error message *vmFail*. The message *vmFail* is also activates the *EnvCleanUp* component, which is responsible for final clean up of the system and the destruction of corresponding VMs. When the clean up is completed, the *EnvCleanUp* component generated the message *scCompleted*, which also indicates that the whole process chain is completed.

If Equation 1 is fulfilled, our platform preforms bootstrapping of generated VMs, and the required compilers are installed according to *ExecParamVM*. If the bootstrapping was successful, the *SCExecution* component is activated by the signal *execStart*. Then we could have two cases:

- The smart connector execution was successful. Then, the *SCExecution* component
  - forwards the results of the computations *dataOutput* to the *OutputTransfer* component;
  - generates the signal *transferStart* that indicates that the *SCExecution* process is completed, and activates *OutputTransfer*.

  The *OutputTransfer* component is responsible for the transfer of the output data to the corresponding server and to a data management system. When the data transfer is completed, the message *transferCompleted* is generated to activate the *EnvCleanUp* component.
- The smart connector execution failed on the stage of scheduling or during execution of a task. Then, the *SCExecution* component generates the signal *execFailed*, to activate *EnvCleanUp* for the final clean up of the system and the destruction of corresponding VMs.

EnvSetUpVM and *EnvCleanUp* can also be logically composed into a meta-component *VMEnv*, which is responsible for any communication with the cloud and the corresponding environment manipulations.

The *SCExecution* component (cf. Figure 3) is the main part of a smart connector. It is responsible for the actual execution of the task and provides a number of the task execution options, defined by parameters *ExecParamT*.

Fault-Tolerance properties of *SCExecution*:
The computation might fail due to network or VM failure, i.e., the VM that hosts some of the processes cannot be reached. To avoid an endless waiting on the output from the processes on the unreachable VM, the smart connector will identify the processes that are hosted there, and then execute the appropriate fault tolerance strategy, e.g., (i) marking the processes that are hosted on the unreachable VM as failed beyond recovery and then collecting the output of processes from the other VMs, or (ii) re-running the failed processes on a different VM until maximum re-run limit is reached. However, we do not implement any strategies to recover a failed process if the failure was due to an internal bug within the task code. In this case, a smart connector will notify the user about the detected failure, as this provides an opportunity to correct the bug.

SCExecution has the following subcomponents:

- *TaskScheduler*: responsible for scheduling of the tasks and their execution in the right order;
- *TaskIterarion*: responsible for execution of task it-

erations according to the corresponding task code; in general case we have NT tasks, where $NT \geq 1$. Thus, a connector has NT components *TaskIterarion*, one for each task (which means that each task should have at least one iteration of its execution);

- *CheckConv* is an optional component, to check whether convergence criterion of a multi-iterational execution is met.

Our model allows us not only to have a precise and concise specification of the cloud-based platform on a logical level but also provides a basis for a formal analysis of its properties, including security properties, as well as of the core computation properties. For the formal analysis we suggest to use an interactive semi-automatic theorem prover Isabelle/HOL (Nipkow et al., 2002; Blanchette et al., 2012) and the corresponding methodologies (Spichkova, 2013b; Spichkova, 2014a; Spichkova, 2014b), as the provided specification is compatible to these methodologies. Moreover, the purposed representation gives a basis for the resource management and performance prediction, cf. (Aversa et al., 2011), as it allows a straightforward analysis of the worst case execution time (WCET) of the composed processes.

The early results on the platform open-source implementation were presented in (Yusuf et al., 2015; Spichkova et al., 2015a). The current version of the platform provides a set of APIs to create new and customise existing SCs. We do not restrict our system to be build using a single programming language. Python was chosen as the development language due to its rapid prototyping features, integration with our data curation system, and due to its increasing uptake by researchers as a scientific software development language. However, the domain-specific calculations could be written in any language. The choice of the language depends on the domain and the concrete research task which should be solved.

3 USABILITY ASPECTS

The proposed open-source platform has been applied across two research disciplines, physics (material characterisation) and structural biology (understanding materials at the atomic scale), to assess its usability and practicality. The domain experts noted the following advantages of the platform:

- time savings for computing and data management,
- user-friendly interface for the computation set up,
- visualisation of the calculation results as 2D or 3D graphs.

The menu has the following sections: *Logout*, *Create Job*, *Jobs*, *Admin*, *Settings*. After logging in, the users are in the *Jobs* section, where they can see the status of current and previous jobs (executions of SCs). Some of the jobs may be processing, have been completed or had errors (cf. Figure 4).

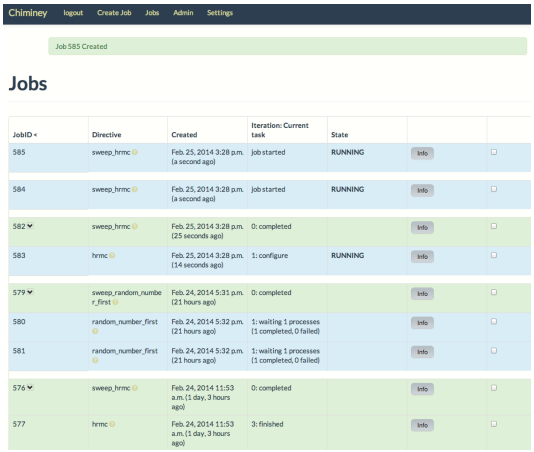

Figure 4: *Jobs* section of the platform.

In the *Settings* section, we can set up general account properties as well as change settings of computation and storage platforms, cf. Figure 5.

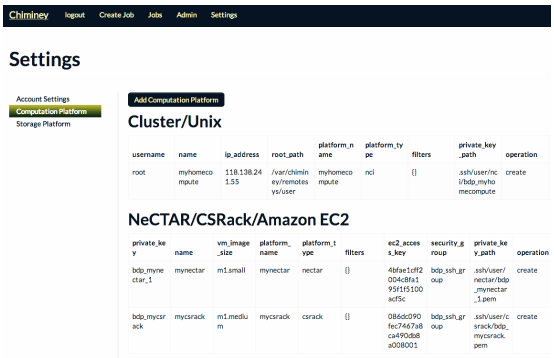

Figure 5: *Settings* section of the platform

When we select *Create Job* in the menu, we will see the job submission page, which has a set of available SCs currently registered. Figure 6 shows how to create a job on example of execution of Monte Carlo simulations, which was a part of one of our case studies (we extended the print screens with the comments to show the match of these parameters to the model from Section 2). In that case study, the Monte Carlo based simulations were applied for modelling of a material's porosity and the size distribution of its pores (industrial applications of these research are in diverse areas such as filtration and gas adsorption). One such

modelling methodology is the Hybrid Reverse Monte Carlo (HRMC), cf. (Opletal et al., 2008). HRMC characterises a material's microstructure by producing models consistent with experimental diffraction data, while at the same time ensuring accurate local bonding environments via energy minimisation.

The user interface, combined with the MyTardis (Androulakis et al., 2008) data curation module, allows for flexible handling of data according to its completion and significance. The results of the calculation can be visualised as 2D or 3D graphs using a plug-in developed to provide better readability of the obtained data (cf. Figures 7 and 8 for examples). The curated datasets are fully accessible and shareable online. The generated graphs can easily be used for presentations or in written documents.

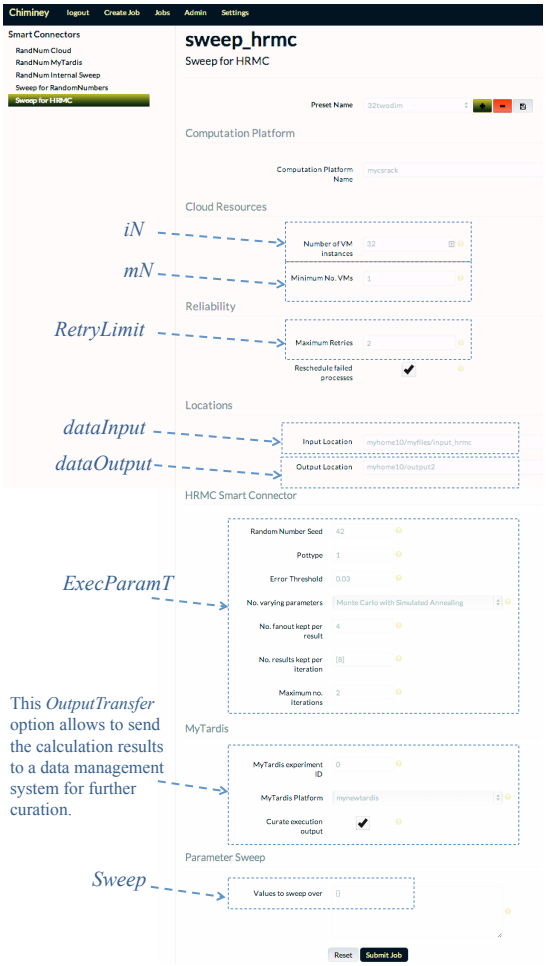

Figure 6: *Create Job* section of the platform.

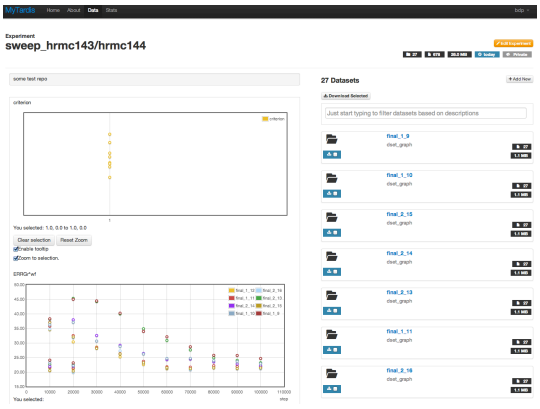

Figure 7: Visualisation of calculation results.

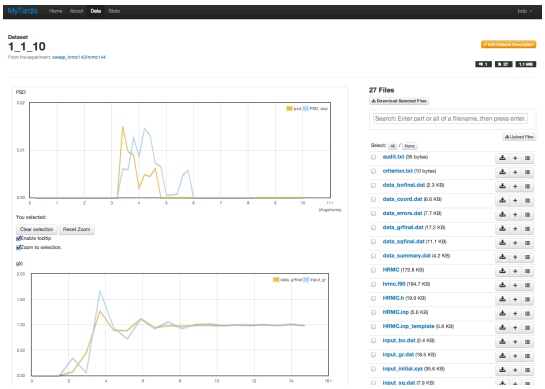

Figure 8: Visualisation of calculation results as 2D graphs.

4 RELATED WORK

While developing the model, we focused on its understandability and readability aspects. There are several approaches on model readability, cf. (Wimmer and Kramler, 2006; Mendling et al., 2007; Zugal et al., 2012; Spichkova et al., 2013; Spichkova et al., 2015b).

The development of formal models and architectures for system involved in cloud computing is a more recent area of system engineering, cf. (Vaquero et al., 2008), (Buyya and Sulistio, 2008), (Leavitt, 2009), (Zhang et al., 2010), (Ostermann et al., 2010), (Cafaro and Aloisio, 2011).

Several approaches have proposed the data stream processing systems for clouds, e.g., (Martinaitis et al., 2009) introduces an approach towards component-based stream processing in clouds, (Kuntschke and Kemper, 2006) presents a work on data stream sharing. Yusuf and Schmidt have shown that the fault-tolerance is best achieved by reflecting the computa-

tional flow in such complex scientific system architectures, cf. (Yusuf and Schmidt, 2013).

There are different types of scientific workflow systems such as Kepler (Ludscher et al., 2006), Taverna (Oinn et al., 2006) and Galaxy (Afgan et al., 2011), which are designed to allow researchers to build their own workflows. The *contribution of the work presented in this paper* is that our platform provides drop-in components, Smart Connectors, for existing workflow engines: (i) researchers can utilise and adapt existing Smart Connectors; (ii) new types of Smart Connectors would be developed within the framework if necessary. From our best knowledge, there is no other framework having this advantage. SCs are geared toward providing power and flexibility over simplicity.

Nimrod (Buyya et al., 2000) is a set of software infrastructure for executing large and complex computational, contains a simple language for describing sweeps over parameter space and the input and output of data for processing. Nimrod is compatible with the Kepler system (Ludscher et al., 2006), such that users can set up complex computational workflows and have them executed without having to interface directly with a high-performance computing system.

One of the directions of our future work is incorporation Nimrod into our open-source platform for the execution of its Smart Connectors. However, Nimrod's web API is currently in development, making interfacing with its capabilities non-trivial in a web-based cloud environment.

5 CONCLUSIONS

Cloud computing provides a great opportunity for scientists, but to unlock all its benefits, we require a platform with a user-friendly interface and an easy-to-use methodology for conducting the experiments. Usability and reliability features are crucial for such systems. This paper presents a model of a cloud-based platform and the latest version of its open-source implementation, focusing on usability and reliability aspects. The proposed platform allows to conduct the experiments without having a deep technical understanding of cloud-computing, HPC, fault tolerance, or data management in order to leverage the benefits of cloud computing.

We believe that the proposed platform will have a strong positive impact on the research community, because it give an opportunity to focus on the main research problems and takes upon itself solving of the major part of the infrastructure problems.

Future work: The main direction of our future work is application of the platform for an efficient testing based on analysis of system architecture.

ACKNOWLEDGEMENTS

The Bioscience Data Platform project acknowledges funding from the NeCTAR project No. 2179 (NeCTAR, 2015).

We also would like to thank our colleagues Dr Daniel W. Drumm (School of Applied Sciences, RMIT University), Dr George Opletal (School of Science, RMIT University), Prof Salvy P. Russo (School of Science, RMIT University), and Prof Ashley M. Buckle (School of Biomedical Sciences, Monash University) for the fruitful collaboration within the Bioscience Data Platform project.

REFERENCES

Afgan, E., Baker, D., Coraor, N., et al. (2011). Harnessing cloud computing with Galaxy Cloud. *Nature Biotechnology*, 29(11):972–974.

Androulakis, S., Schmidberger, J., Bate, M. A., et al. (2008). Federated repositories of X-ray diffraction images. *Acta Crystallographica, Section D*, 64(7):810–814.

Armbrust, M., Fox, A., Griffith, R., et al. (2010). A view of cloud computing. *Commun. ACM*, 53(4):50–58.

Aversa, R., Di Martino, B., Rak, M., Venticinque, S., and Villano, U. (2011). *Performance Prediction for HPC on Clouds*, pages 437–456. John Wiley & Sons, Inc.

Blanchette, J., Popescu, A., Wand, D., and Weidenbach, C. (2012). More SPASS with Isabelle – Superposition with hard sorts and configurable simplification. In Beringer, L. and Felty, A., editors, *Interactive Theorem Proving*, volume 7406 of *LNCS*, pages 345–360. Springer Berlin Heidelberg.

Buyya, R., Abramson, D., and Giddy, J. (2000). Nimrod/G: An Architecture for a Resource Management and Scheduling System in a Global Computational Grid.

Buyya, R. and Sulistio, A. (2008). Service and utility oriented distributed computing systems: Challenges and opportunities for modeling and simulation communities. In *Proc. of the 41st Annual Simulation Symposium*, ANSS-41 '08, pages 68–81. IEEE Comp. Society.

Buyya, R., Yeo, C. S., Venugopal, S., Broberg, J., and Brandic, I. (2009). Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. *Future Generation Computer Systems*, 25(6):599 – 616.

Cafaro, M. and Aloisio, G. (2011). Grids, clouds, and virtualization. In Cafaro, M. and Aloisio, G., editors, *Grids, Clouds and Virtualization*, Computer Communications and Networks, pages 1–21. Springer London.

Kuntschke, R. and Kemper, A. (2006). Data stream sharing. In Grust, T., Höpfner, H., Illarramendi, A., Jablonski, S., Mesiti, M., Müller, S., Patranjan, P.-L., Sattler, K.-U., Spiliopoulou, M., and Wijsen, J., editors, *Current Trends in Database Technology – EDBT 2006*, volume 4254 of *LNCS*, pages 769–788. Springer.

Leavitt, N. (2009). Is cloud computing really ready for prime time? *Computer*, 42(1):15–20.

Ludscher, B., Altintas, I., Berkley, C., et al. (2006). Scientific workflow management and the Kepler system. *Concurrency and Computation: Practice and Experience*, 18(10):1039–1065.

Martinaitis, P. N., Patten, C. J., and Wendelborn, A. L. (2009). Component-based stream processing "in the cloud". In *Proceedings of the 2009 Workshop on Component-Based High Performance Computing*, CBHPC '09, pages 16:1–16:12. ACM.

Mendling, J., Reijers, H., and Cardoso, J. (2007). What makes process models understandable? In Alonso, G., Dadam, P., and Rosemann, M., editors, *Business Process Management*, volume 4714 of *LNCS*, pages 48–63. Springer.

NeCTAR (2015). The National eResearch Collaboration Tools and Resources. `http://www.nectar.org.au/`.

Nipkow, T., Paulson, L. C., and Wenzel, M. (2002). *Isabelle/HOL – A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer.

Oinn, T., Greenwood, M., Addis, M., et al. (2006). Taverna: Lessons in Creating a Workflow Environment for the Life Sciences. *Concurr. Comput. : Pract. Exper.*, 18:1067–1100.

Opletal, G. et al. (2008). HRMC: Hybrid Reverse Monte Carlo method with silicon and carbon potentials. *Computer Physics Communications*, (178):777–787.

Ostermann, S., Iosup, A., Yigitbasi, N., Prodan, R., Fahringer, T., and Epema, D. (2010). A Performance Analysis of EC2 Cloud Computing Services for Scientific Computing. In Avresky, D., Diaz, M., Bode, A., Ciciani, B., and Dekel, E., editors, *Cloud Computing*, volume 34 of *Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering*, pages 115–131. Springer Berlin Heidelberg.

Spichkova, M. (2011). Focus on processes. Tech. Report TUM-I1115, TU München.

Spichkova, M. (2013a). Design of formal languages and interfaces: "formal" does not mean "unreadable". In Blashki, K. and Isaias, P., editors, *Emerging Research and Trends in Interactivity and the Human-Computer Interface*. IGI Global.

Spichkova, M. (2013b). Stream Processing Components: Isabelle/HOL Formalisation and Case Studies. *Archive of Formal Proofs*.

Spichkova, M. (2014a). Compositional properties of crypto-based components. *Archive of Formal Proofs*.

Spichkova, M. (2014b). Formalisation and analysis of component dependencies. *Archive of Formal Proofs*.

Spichkova, M. and Schmidt, H. (2015). Reconciling a component and process view. *7th International Workshop on Modeling in Software Engineering (MiSE) at ICSE 2015*.

Spichkova, M., Thomas, I., Schmidt, H., Yusuf, I., Drumm, D., Androulakis, S., Opletal, G., and Russo, S. (2015a). Scalable and fault-tolerant cloud computations: Modelling and implementation. In *Proc. of the 21st IEEE International Conference on Parallel and Distributed Systems (ICPADS 2015)*.

Spichkova, M., Zamansky, A., and Farchi, E. (2015b). Towards a human-centred approach in modelling and testing of cyber-physical systems. In *Proc. of the International Workshop on Automated Testing for Cyber-Physical Systems in the Cloud*.

Spichkova, M., Zhu, X., and Mou, D. (2013). Do we really need to write documentation for a system? In *International Conference on Model-Driven Engineering and Software Development (MODELSWARD'13)*.

Vaquero, L. M., Rodero-Merino, L., Caceres, J., and Lindner, M. (2008). A break in the clouds: towards a cloud definition. *SIGCOMM Comput. Commun. Rev.*, 39(1):50–55.

Wimmer, M. and Kramler, G. (2006). Bridging grammarware and modelware. In Bruel, J.-M., editor, *Satellite Events at the MoDELS 2005 Conference*, volume 3844 of *LNCS*, pages 159–168. Springer.

Yusuf, I. and Schmidt, H. (2013). Parameterised architectural patterns for providing cloud service fault tolerance with accurate costings. In *Proc. of the 16th Int. ACM Sigsoft symp. on Component-based software engineering*, pages 121–130.

Yusuf, I., Thomas, I., Spichkova, M., Androulakis, S., Meyer, G., Drumm, D., Opletal, G., Russo, S., Buckle, A., and Schmidt, H. (2015). Chiminey: Reliable computing and data management platform in the cloud. In *Proc. of the International Conference on Software Engineering (ICSE'15)*, pages 677–680.

Zhang, Q., Cheng, L., and Boutaba, R. (2010). Cloud computing: state-of-the-art and research challenges. *Journal of Internet Services and Applications*, 1(1):7–18.

Zugal, S., Pinggera, J., Weber, B., Mendling, J., and Reijers, H. (2012). Assessing the impact of hierarchy on model understandability – a cognitive perspective. In Kienzle, J., editor, *Models in Software Engineering*, volume 7167 of *Lecture Notes in Computer Science*, pages 123–133. Springer.